

UNCLASSIFIED

**Defense Technical Information Center
Compilation Part Notice**

ADP019943

TITLE: On the Road Towards Superconductor Computers: Twenty Years Later

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Future Trends in Microelectronics: The Nano, the Giga, and the Ultra. IEEE International Conference on Nanotechnology [4th] Held in Munich, Germany on 16-19 August 2004

To order the complete compilation report, use: ADA436172

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP019925 thru ADP019959

UNCLASSIFIED

On the Road Towards Superconductor Computers: Twenty Years Later

M. Dorojevets, D. Strukov

*Dept. of Electrical and Computer Engineering
SUNY-Stony Brook, NY 11794, USA*

A. Silver, A. Kleinsasser

Jet Propulsion Laboratory (NASA), Pasadena, CA, USA

F. Bedard

Department of Defense, USA

P. Bunyk, Q. Herr, G. Kerber, and L. Abelson

Northrop Grumman Space Technology, Redondo Beach, CA USA

1. Introduction

Attempts to build superconductor computers based on low-temperature superconductivity of Josephson junctions have come and gone during the last thirty years. In spite of both industry and the federal government's deep involvement in these efforts, no full-fledged superconductor computers have been built so far. As a result, superconductors, once considered an alternative to semiconductors, have lost their appeal to the general design community. By 2003, tremendous improvements in silicon-based technology have allowed semiconductor chips to have hundreds of millions of transistors and clock frequencies exceeding 3 GHz, speeds earlier considered the province of exotic technologies, such as superconductors.

Although superconductors have not become a mainstream digital technology (and perhaps it is fair to say they were never expected to be), they are still trying to find their way into several important niche areas, one of which is extreme (at the moment called petaflops) computing. Thanks to the support by various U.S. government agencies, work on superconductor technology and processor design continued behind the scenes, nearly invisible to silicon-based computer designers. In this paper, we consider the lessons learned from past superconductor computer projects, the current status, and future directions of the work in this field.

2. IBM Josephson computer technology project

The first full-scale attempt to implement superconductor computing was a program at IBM in the 1972–1983 time frame, using the first-generation superconductor

technology based on Josephson junction (JJ) devices in "latching" logic circuits. The goal was to develop the technology to a level where a significant, useful high performance computer would be viable. The potential customers included both "the corner bank" and the special users requiring the highest computation rates. In the device domain, fabrication processes, a circuit family, packaging and powering techniques were developed and demonstrated. As a demonstration of system performance, a digital signal processor called JSP (for Josephson signal processor) was designed, based upon a semiconductor version. The project then focused on building and testing this JSP until the program was terminated in September of 1983. In retrospect, the goals of this project were clock frequencies about 1 GHz, the highest that could be expected with the fabrication and circuit technology available at that time. The IBM fabrication technology was improved upon during the MITI project in Japan in the 1980's, particularly with the change from niobium-lead alloy junctions and lead-alloy wiring to all-refractory niobium-based fabrication. Fujitsu, NEC, Hitachi, as well as Japan's ETL were funded under this project, but despite small processor and memory demonstrations, no full-fledged supercomputing effort was mounted.

3. Hybrid technology multi-threaded petaflops computer project

The hybrid technology multi-threaded (HTMT) architecture petaflops project (1996-2000)¹ presented another chance to resurrect superconductor computing after almost 15 years since the end of the IBM-led project in US and the MITI project in Japan.

New hopes for success were based on several breakthroughs in circuit design such as development a new rapid single flux quantum (RSFQ)² logic using the magnetic flux quantization properties of superconducting rings with Josephson junctions, as well as the development of niobium-trilayer technological processes in the U.S. This RSFQ logic represented a significant improvement over previous superconductor logic families employed by IBM and the MITI program in Japan. Those earlier programs were based on voltage-latching gates that dissipated more power and were limited to several GHz by the need to actively reset each gate every clock cycle. In addition, latching logic required multi-phase rf power systems, while RSFQ logic uses only dc power.

RSFQ circuits operate by the creation, elimination, and propagation of single magnetic flux quanta ($\phi_0 = h/2e = 2.07 \text{ mV}\cdot\text{ps}$) in small monolithic thin-film inductive loops (a few picohenries) containing Josephson tunnel junctions. Josephson junctions can switch and perform logic functions in a few picoseconds. Bits are stored as a persistent supercurrent in the inductor, at zero voltage. Bits propagate from one gate to the next as millivolt picosecond SFQ pulses, not voltage levels. Matched strip and microstrip lines are used to propagate signals beyond the adjacent gate.

These technological improvements coincided with the decision by policy makers at U.S. federal agencies that computer processing capabilities reaching and exceeding one petaflop (10^{15} floating-point operations per second) were in the national interest. It was recognized that massively parallel systems designed with

thousands of silicon-based processors could reach this target not earlier than 2010. As suggested by the authors of the HTMT concept and later confirmed by the results of the HTMT studies, with adequate investments into technology and design tools, superconductor processors coupled with other new technologies and architectural concepts could provide a "shortcut" into this petaflops computing territory. Thus, in 1996 superconductor processor design got a new life with a new focus on extreme, petaflops level computing, where it could outperform, in terms of speed and power, semiconductor processors.

The results of these HTMT design studies related to the area of superconductor computing have been the following:

- the development of a parallel superconductor processor element (Spell) architecture³ with two-level simultaneous multithreading, capable of addressing the huge disparities in cycle times and access latency between superconductor Spell processors and silicon-based memory;
- the feasibility analysis of possible implementation of a multiprocessor superconductor sub-subsystem⁴ consisting of 4,096 ~50 GHz Spell processors with their local small superconductor memory interconnected by a multi-stage superconductor switch, all implemented in the (still to be developed) $0.8\text{ }\mu\text{m}$, 20 kA/cm^2 superconductor technology;
- efficient integration of the superconductor sub-system with other components of a hierarchical HTMT system;
- development of new 4 kA/cm^2 critical current density, $1.75\text{ }\mu\text{m}$ Nb/Al-AlO_x/Nb trilayer superconductor process⁵ at the former TRW, Inc. (now Northrop Grumman Space Technology – NGST). Since that time, NGST has advanced to an 8 kA/cm^2 , $1.25\text{ }\mu\text{m}$ process and is collaborating with JPL to initiate a 20 kA/cm^2 , $0.8\text{ }\mu\text{m}$ process.

In parallel with these developments at SUNY-Stony Brook and TRW, Prof. Van Duzer and his team at UC Berkeley worked on an RSFQ chip-level interface for a 64-Kbit CMOS RAM with a projected read access time of ~600 ps at 4 °K.⁶ The relatively high power dissipation for these Josephson-CMOS hybrid memory chips remains a concern, so the lack of fast, dense, and low power superconductor memory remains one of the longstanding problems to be overcome before petaflops computing with superconductor processors becomes feasible.

Outside the HTMT project, another project known as the Superconductive Crossbar was initiated by the U.S. Dept. of Defense in 1990's, to demonstrate the feasibility of using superconductor technology for very high data rate communication among processors as well as shared memory in parallel multiprocessor systems. The architecture for such a switch was a 128x128 self-routing crossbar capable of processing data streams at a rate of 2.5 Gbit per second per channel.

Although the results of the four years of the HTMT studies were optimistic about the potential use of superconductor technology for supercomputing, they also highlighted the huge gap between the current state of superconductor circuit design and VLSI fabrication, and the high complexity and reliability required from

superconductor chips for such a petaflops system. By 2000, the most complex superconductor LSI chips had only a few thousands junctions, capable of doing only some non-programmable, built-in signal processing functions. By comparison, the HTMT Spell processors were expected to have up to 400K gates, or several million junctions per chip. Although other new HTMT technologies were also in experimental stages of development, none of them created more controversy than superconductor processors among reviewers of the project. It was perhaps not surprising that sponsors were reluctant to commit several hundred million dollars for the full development, design, and fabrication of an HTMT prototype system with several thousands of superconductor processors until at least a small superconductor microprocessor prototype had been demonstrated.

4. 8-bit FLUX-1 microprocessor

In order to address these issues and get a practical handle on real-world superconductor VLSI microprocessor chip design, the collaboration between SUNY Stony Brook, TRW, and JPL (NASA) to design and demonstrate a FLUX-1 microprocessor chip started in 2000.⁷ A FLUX-1 chip was expected to be a superconductor technology "driver", without any plans for using it in future superconductor systems. Its design was based on NGST's new 4 kA/cm² 1.75 μm Nb-trilayer superconductor fabrication process.⁵

Figure 1 shows the block diagram of the 20 GHz 8-bit FLUX-1 microprocessor, the most ambitious chip implemented in RSFQ technology to date. It includes a small instruction memory of 16 30-bit instructions with the embedded program counter and instruction fetch logic, the branch unit, the instruction register and dual decode/issue logic, eight bit-stream arithmetic logic units (ALUs) interleaved with eight 8-bit general-purpose integer registers (REG0–REG7), two 8-bit 20 GHz I/O ports, the clock controller, and built-in scan path circuitry. FLUX-1 has a partitioned deeply-pipelined, synchronous dual-op long-instruction-

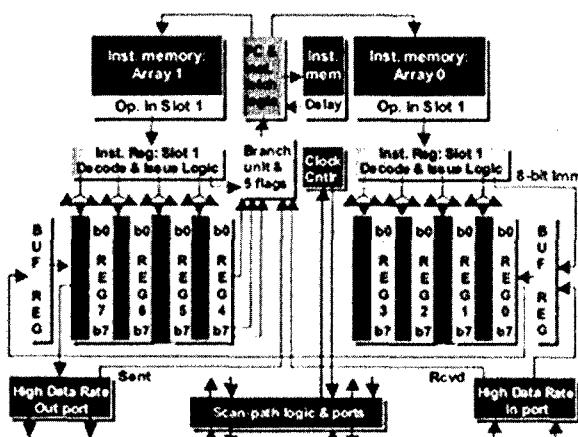


Figure 1. Block diagram of a FLUX-1 microprocessor.

word architecture with an instruction set of ~25 control, integer arithmetic, and logical operations. Two instructions can be issued and completed during each 50 ps cycle.⁸⁻⁹

The functional correctness of the architecture was verified with a cycle-accurate FLUX-1 simulation model. During the physical implementation of the chip, the chip layout was done manually due to the lack of synthesis and automatic placement and routing tools tuned for superconductor technology.

The first wafers with FLUX-1 chips were delivered in August of 2001. Because of a very high 6.9 A bias current supplied to these chips, the power pads on the FLUX-1 chip carrier had to be redesigned to avoid overheating. The large chip size of 140 mm² made it very difficult to find dies without any defects. Nevertheless, the most serious problem was the very high bit error rate (BER) in the FLUX-1 scan path circuits, which made further chip testing impossible. It was concluded that this high BER was probably due to errors in the circuit-level design of transmission line drivers and receivers. As a result, both the FLUX-1 gate library and the chip layout were improved. First wafers of the revised FLUX-1R chip, with the same architecture but different circuit-level design, were fabricated in the summer of 2002.¹⁰ FLUX-1R was built using a library of only 10 SFQ gates, which were designed to be interconnected with passive matched transmission lines.¹¹⁻¹² A FLUX-1R chip shown in Fig. 2 has 21% less chip area, 33% less bias current, 4% fewer Josephson junctions, and other improvements over the first version of the chip. The total power dissipation is ~10 mW.

The complete FLUX-1R stripline-connected gate library has been verified and the correct operation of several FLUX-1R microprocessor blocks has been demonstrated, including the 13-stage scan path within the first instruction decoder on the FLUX-1R chip itself, a significant advance over the first design. Measured low-frequency gate bias margins (see Table 1) were large, reproducible between chips from different wafers, and in good agreement with simulations. Equally important, all gate bias margins strongly overlap.

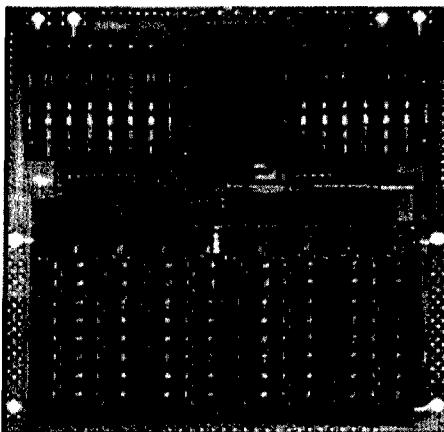


Figure 2. FLUX-1R chip photomicrograph. There are 63,107 Josephson junctions on a 10.35x10.65 mm² die.

A number of larger breakout structures were also tested successfully, although operating margins were smaller. To date, the largest circuit block successfully tested with reproducible margins across different chips is a 1-bit ALU register block that is a key element of the FLUX-1 computational engine, where this block is replicated 64 times to create a parallel bit-stream data processing data path. This circuit is a microcosm of the FLUX-1R chip and makes liberal use of passive interconnects, including crossovers. Further testing of the FLUX-1R chip is planned.^{10,13}

Among other results of the FLUX-1 project, the most important ones are:

- the successful demonstration of inter-chip, flux-quantum, serial-data transmission up to 60 Gb/s through a passive substrate,¹⁴ and
- development of a new advanced 8 kA/cm², 1.2 μm junction fabrication process,¹⁵ as the next step in the evolution towards higher speed and increased circuit density.

The new process is based on NGST's 4 kA/cm² Nb process used previously for the FLUX-1R chip fabrication, but several new process steps have been developed to improve yield and reproducibility. Minimum junction size and line-pitch in the 8 kA/cm² process are 1.2 μm and 2.6 μm, respectively. Critical current spreads are typically 1.5% (1σ) for the arrays of 1.2 μm junctions, comparable to the best spreads achieved in the 4 kA/cm² process. The new process has been used to demonstrate a 300 GHz static digital divider and to fabricate complex digital circuits of 28,000 junctions.

Gate	V bias (mV)	Margin (mV)	Margin (%)
XOR	1.74	0.75	43%
SPLIT	1.63	0.70	43%
DFF	1.79	0.70	39%
INV	1.63	0.60	37%
MERGER	1.79	0.65	36%
D2FF	1.68	0.55	33%
INPUT	1.75	0.55	31%
AND	1.79	0.50	28%
OUTPUT	1.82	0.50	27%
NDRO	1.73	0.45	26%
All gates	1.73	0.45	26%

Table 1. FLUX-1R gate library bias margins.

5. FLUX-2: 32-bit floating-point vector multiplier chipset

In 2002, in parallel with the testing a FLUX-1R microprocessor chip at NGST, work started on a new FLUX-2 project, a 32-bit floating-point (FP) vector multiplier unit chipset with a target clock frequency of 25 GHz, by the same team of SUNY-Stony Brook, NGST, and JPL. The primary goal of the FLUX-2 project is to shed light on how complex processing units operating at 25+ GHz clock frequencies and consisting of multiple superconductor chips can be designed, fabricated, and packaged on a single MCM carrier. This is a crucial step for superconductor computer design because the low gate density of the current, and perhaps even next generation, superconductor chips will not allow a 32/64-bit processor to be implemented on a single chip.

Below we discuss in detail the major results of the work in progress done by the SUNY-Stony Brook team that has been developing the architecture and gate-level design of a dual-chip 32-bit FP vector multiplier of the FLUX-2 chipset.

Our goal is to design a dual-chip multiplier capable of calculating one 32-bit floating-point result per 40 ps cycle, while processing two input streams of data encoded in the IEEE 754 single precision format from vector register memory on a separate chip. All data transfers between chips placed on a multi-chip module are to be done over superconductor Nb transmission lines at the same clock rate used to process and read/write data inside the chips.

The key problems to be solved are the following:

- partitioning of the 32-bit FP pipelined vector multiply unit into multiple chips of $\sim 1 \text{ cm}^2$ size each, with a gate count ($\sim 10\text{K}$) and the number of I/O pads of each chip within the capabilities of NGST's $1.2 \mu\text{m}$ chip fabrication and solder reflow packaging processes;
- tolerating significant clock and data skew both between the chips on an MCM carrier and between blocks within each chip;
- processing of vector data from memory at the 25 GHz rate.

The top-level functional block diagram for a FLUX-2 floating point multiplier is shown in Fig. 3. A 32-bit FP multiplier consists of an unsigned 24-bit integer multiplier for significands (mantissas), 8-bit adder for exponents, sign calculation by XORing the signs of two operands, plus support circuitry to deal with the exponents and special values (0, +inf, -inf, etc.). The significand is in the [1, 2] range, so the product falls in the [1, 4] range. Thus, normalizing as well as incrementing a tentative exponent may be required. The IEEE floating point multiplier has four different rounding schemes,^{16,17} necessitating a second set of "normalize" and "adjusting exponent" steps after the "round" step,¹⁸ because for some of the schemes the result after rounding needs to be shifted. In our case, we implemented only truncation rounding which does not require that extra step.

The slowest and the most challenging part of a multiplier is multiplication of significands. Indeed, the largest part of the critical path delay (more than 90% in our case) comes from the 24-bit unsigned integer multiplication of significands.

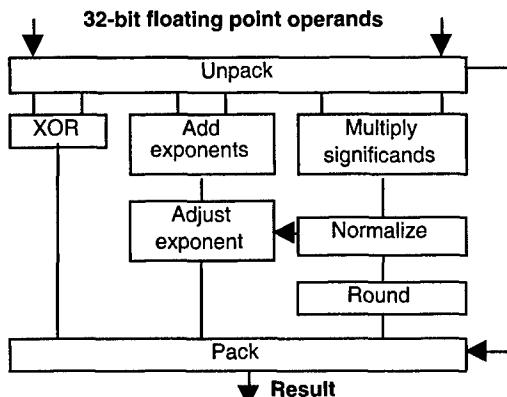


Figure 3. Functional block diagram of a FLUX-2 floating-point multiplier.

Many different algorithms have been proposed for fast integer multiplication.²⁰ Logically, it can be separated into three stages: partial product (PP) generation, PP reduction, and final summation.

While Booth's algorithm for PP generation together with the Wallace/binary tree for the PP reduction can be found in most CMOS high-speed multipliers,^{20,21} this combination does not provide the best solution in the case of an RSFQ multiplier because the cost of broadcasting signals with RSFQ splitters is high (in terms of latency and data skew). Further, the relatively small number of metal layers available with the present superconductor process makes it very difficult to lay out the Wallace/binary trees without wasting too much valuable chip area.

For the case of 24-bit multiplication implemented in RSFQ logic, our analysis of the Booth 2 encoding with different PP reduction topologies showed that the use of Booth 2 can give only insignificant performance gains, decreasing the multiply latency by one cycle at most for some of the topologies. In the meantime, the savings in area due to the decreased number of PP's with the Booth 2 encoding (13 with Booth 2 vs. 24 without it) are compensated by the increase in area to lay out the significantly larger number of wires required to broadcast signals. This analysis precluded the use of any Booth encoding in the PP generation stage.

In theory, a binary tree used during the PP reduction phase of multiplication of significands gives the best performance in terms of the number of pipeline stages, but performance degrades when it is implemented in RSFQ logic. The superior performance of a binary tree (or any tree topology in general) comes at a heavy price. Uneven concentration of wires in different pipeline stages complicates the physical layout and produces wasted area. The less compact layout adds additional pipeline stages in the critical path. We chose the (4-4-6-10) high-order array topology for the PP reduction because it is easier to lay out, and it gives the same performance as the binary tree.

The advantage of a high-order array is that the number of connections between sub-arrays is constant. A two-dimensional structure of such array is easy to lay out, especially given the fact that all rows are about of the same width, i.e., have

the same number of compressors. The layout is very regular and implemented by tiling of modules. There are total of 8 different module types, each with three [4:2] compressors and some with the PP encoding logic and signal propagation latches. By breaking long wires between rows (which combine results from sub-arrays) and aligning rows in a specific manner, we developed a systolic-like structure where each module is only connected to its neighbors. The modules are connected in a corner-based clock network. A clock distribution network with very small clock skew is implemented inside each module.

Truncation rounding is performed in parallel with the reduction and the encoding stages by adding the least significant bits coming from the high-order array in a ripple-carry fashion. The sum bits are dropped, while a carry value is propagated to the final adder. In the final stage, the Kogge-Stone algorithm²² is executed to add two 32-bit numbers. The exponent adder is not on the critical path, whereas the exponent adjustment is on it. In order to reduce overall latency, adjusting of an exponent is implemented by calculating two sums simultaneously, the sum of two exponents and the sum of two exponents plus one, with each sum calculated by a separate ripple-carry adder. If the result from the final adder has to be normalized, then the second sum is chosen to be a new exponent value. The first sum is selected otherwise.

Table 2 shows the estimated latency, area, and complexity for different blocks of the FLUX-2 dual chip multiplier. All physical characteristics of the chips were estimated for the gate-level multiplier schematics developed at SUNY-Stony Brook, using NGST's j110D FLUX-1 gate library, with gates as well as the wire pitch scaled down to those in the NGST's new j110E 8 kA/cm², 1.2 μm process.

	FP multiplier chip 1		FP Multiplier chip 2			Dual-chip 32-bit FP multiplier total
	Encoding and PP reduction	Rounding	Final adder	Exponent adder	Normalize significand and adjust exponent	
Latency, 40 ps cycles	18	not on the critical path	8–10	not on the critical path	1	27–29 (w/o chip- to-chip wire delay)
JJ count	~105K	~1.2K	~17K	~8.3K	~1.5K	~135K
Area, mm ²	~ 6.5 x 6		~ 3 x 3			~ 10 x 10

Table 2. Estimated design parameters for a 25 GHz 32-bit FP dual-chip multiplier.

6. Conclusions

Despite the progress that has been made, important technology challenges remain to be addressed before the first practical superconductor processor can be demonstrated, namely:

- Nb VLSI fabrication process needs further improvements to provide a higher clock frequency, greater gate density, and greater reliability;
- At least 64 Kb, 0.1 ns low-power cryogenic RAM chips have to be developed;
- Circuit designs that enable lower power dissipation, lower bias current per chip, and higher clock rates are needed;
- Reliable MCM packaging and parallel data chip-to-chip communication techniques are to be demonstrated;
- Thermally and electrically efficient, high-data-rate cryogenic-to-ambient I/O technology has to be developed;
- Efficient CAD tools to design and verify RSFQ circuits operating at 25+ GHz clock frequencies are required.

In the architecture domain, the realities of the current and future Nb fabrication processes (especially the limited gate density of Nb VLSI logic and memory chips) and the peculiarities of the RSFQ logic must be fully addressed in the design of RSFQ processors and its blocks in order to utilize the huge potential of superconductors for petaflops computing.

Acknowledgments

This work was supported by U.S. Dept. of Defense.

References

1. G. Gao, K. Likharev, P. Messina, and T. Sterling, "Hybrid technology multithreaded architecture," in: *Proc. 6th Symp. Frontiers Massively Parallel Computation*, Los Alamitos, CA: IEEE Computer Society Press, 1996, pp. 98–105.
2. K. Likharev and V. Semenov, "RSFQ logic/memory family: a new Josephson junction technology for sub-terahertz clock frequency digital systems," *IEEE Trans. Appl. Supercond.* **1**, 3 (1991).
3. M. Dorojevets, "COOL multithreading in HTMT SPELL-1 processors," in Y. S. Park, S. Luryi, M. S. Shur, J. M. Xu, and A. Zaslavsky, eds., *Frontiers in Electronics: From Materials to Systems*, Singapore: World Scientific, 2000, pp. 247–253.
4. M. Dorojevets, P. Bunyk, D. Zinoviev, and K. Likharev, "RSFQ computing: The quest for petaflops," in: S. Luryi, J. M. Xu, and A. Zaslavsky, eds., *Future Trends in Microelectronics: The Road Ahead*, New York: Wiley, 1999, pp. 193–206.
5. G. L. Kerber, L. A. Abelson, M. L. Leung, Q. P. Herr, and M. W. Johnson, "A high density 4 kA/cm^2 Nb integrated circuit process," *IEEE Trans. Appl. Supercond.* **9**, 1061 (2001).
6. Y. J. Feng, X. Meng, S. R. Whiteley, T. Van Duzer, K. Fujiwara, H. Miyakawa, and N. Yoshikawa, "Josephson-CMOS hybrid memory with ultra-high-speed interface circuit," *IEEE Trans. Appl. Supercond.* **13**, 467 (2003).
7. M. Dorojevets, P. Bunyk, and D. Zinoviev, "FLUX chip: Design of a 20-GHz 16-bit ultrapi pipelined RSFQ processor prototype based on 1.75- μm LTS technology," *IEEE Trans. Appl. Supercond.* **11**, 326 (2001).
8. M. Dorojevets, "An 8-bit FLUX-1 RSFQ microprocessor built in 1.75- μm technology," *Physica C* **378-381**, 1446 (2002).
9. M. Dorojevets and P. Bunyk, "Architectural and implementation challenges in designing high-performance RSFQ processors: FLUX-1 microprocessor and beyond," *IEEE Trans. Appl. Supercond.* **13**, 446 (2003).
10. P. Bunyk, M. Leung, J. Spargo, and M. Dorojevets, "FLUX-1 RSFQ microprocessor: Physical design and test results," *IEEE Trans. Appl. Supercond.* **13**, 433 (2003).
11. D. Zinoviev, P. Bunyk, and P. Konig, "Passive interconnect: A revolutionary approach to RSFQ system design," *8th Intern. Supercond. Electronics Conf.* Osaka, Japan, June 2001.
12. Q. Herr, M. Wire, and A. Smith, "Ballistic SFQ signal propagation on-chip and chip-to-chip," *IEEE Trans. Appl. Supercond.* **13**, 463 (2003).
13. L. Abelson, P. Bunyk, M. Dorojevets, Q. Herr, G. Kerber, A. Kleinsasser, A. Silver, and D. Strukov, "Development of superconductor electronics technology for high end computing," submitted to *J. Supercond. Sci. Technol.* (2003).
14. Q. P. Herr, A. D. Smith, and M. S. Wire, "High-speed data link between digital superconducting chips," *Appl. Phys. Lett.* **80**, 3210 (2002).

15. G. L. Kerber, L. A. Abelson, K. Edwards, R. Hu, M. W. Johnson, M. L. Leung, and J. Luine, "Fabrication of high current density Nb integrated circuits using a self-aligned junction anodization process," *IEEE Trans. Appl. Supercond.* **13**, 82 (2003).
16. ANSI/IEEE 754-1985: "IEEE standard for binary floating-point arithmetic," also in *Computer* **14**, 51 (1981).
17. Israel Koren, *Computer Arithmetic Algorithms*, Natick, MA: A. K. Peters Ltd., 2001.
18. Behrooz Parhami, *Computer Arithmetic: Algorithms and Hardware Designs*, Oxford: Oxford University Press, 1999.
19. Michael J. Flynn and Stuart F. Oberman, *Advanced Computer Arithmetic Designs*, New York: Wiley, 2001.
20. G. Hinton, D. Sager, M. Upton, D. Boggs, D. Carmean, A. Kyker, and P. Roussel, "The microarchitecture of the Pentium 4 processor," *Intel Technol. J.* Q1 (2001).
21. E. M. Schwarz and C. A. Krygowski, "The S/390 G5 floating point unit," *IBM J. Res. Develop.* **43**, 707 (1999).
22. P. M. Kogge and H. S. Stone, "A parallel algorithm for the efficient solution of a general class of recurrence equations," *IEEE Trans. Comp.* **22**, 786 (1973).